When and Why the Topological Coverage Criterion Works[*]

Nicholas J. Cavanna[†]   Kirk P. Gardner[‡]   Donald R. Sheehy[§]

Abstract

In their seminal work on homological sensor networks, de Silva and Ghrist showed the surprising fact that it's possible to certify the coverage of a coordinate-free sensor network even with very minimal knowledge of the space to be covered. Here, coverage means that every point in the domain (except possibly those very near the boundary) has a nearby sensor. More generally, their algorithm takes a pair of nested neighborhood graphs along with a labeling of vertices as either boundary or interior and computes the relative homology of a simplicial complex induced by the graphs. This approach, called the Topological Coverage Criterion (TCC), requires some assumptions about the underlying geometric domain as well as some assumptions about the relationship of the input graphs to the domain. The goal of this paper is to generalize these assumptions and show how the TCC can be applied to both much more general domains as well as very weak assumptions on the input. We give a new, simpler proof of the de Silva-Ghrist Topological Coverage Criterion that eliminates any assumptions about the smoothness of the boundary of the underlying space, allowing the results to be applied to much more general problems. The new proof factors the geometric, topological, and combinatorial aspects, allowing us to provide a coverage condition that supports thick boundaries, k-coverage, and weighted coverage, in which sensors have varying radii.

1 From Sensor Coverage to Data Coverage

Problems in Homological Sensor Networks (HSNs) are usually stated in the vocabulary of sensor networks. There are sensors and coverage regions, and an important problem is to determine if the sensing region of a collection of sensors covers a given domain, given only the neighborhood relationships between the sensors and some indication of which sensors are near the boundary. The locations (coordinates) of the points is not assumed, nor is the shape (topology) of the domain to be covered. Although phrased in the language of sensor networks, the problem may be understood more generally as one of data coverage; it answers when a data set

sufficiently covers a domain, and holes in coverage can be viewed as gaps in the data.

A surprising result by de Silva & Ghrist is that there exists a computable, sufficient condition called the *Topological Coverage Criterion (TCC)* to certify coverage without knowing the locations of the sensors when the domain's boundary is smooth [4]. All that is required is that the sensors have unique identifiers, can detect nearby sensors, can differentiate whether neighboring sensors are "close" or "very close" (in a technical sense to be defined below), and can detect if the boundary of the domain is close. Developed over a series of papers [8, 17, 5, 4], the most general version of the TCC can be understood as a purely geometric problem on point sets P in $\mathbb{R}^d$ with unknown coordinates in an unknown domain $\mathcal{D}$. The input is just a pair of nested neighborhood graphs on the points indicating which points are close to which others and a labeling of the points indicating which points are close to the boundary of $\mathcal{D}$. Determining coverage is ill-posed with only this information, but de Silva & Ghrist proved that the TCC gives a sufficient condition for coverage (of a slightly shrunken version) of the domain when the radii are all equal, the boundary of the domain is a smooth manifold with bounded injectivity radius, and the goal is to certify single-coverage (each point covered by at least one sensor). *In this paper, we generalize these results in several directions, most notably in eliminating many of the geometric assumptions on the domain, while also certifying multiple-coverage and allowing for varying radii.*

For an input graph $G = (P, E)$ where P is a set of points, and a subset $Q \subseteq P$ of points labeled as near the boundary, the TCC algorithm builds a pair of simplicial complexes, one from the cliques in G and one from the cliques in the subgraph of G induced by Q. By a straightforward matrix reduction, a vector space called the relative d-dimensional homology of this nested pair of subcomplexes can be computed. In an ideal situation, the dimension would indicate exactly the number of connected components if and only if the domain is covered. However, this simplistic approach does not work in general, because it can detect spurious components. The insight of de Silva and Ghrist was to show that for sufficiently nice geometric

domains, such as subsets of $\mathbb{R}^d$ bounded by sufficiently smooth manifolds, the spurious components can be eliminated by using two different scales and replacing the homology with the so-called persistent homology induced by inclusion of the smaller pair of complexes into the larger pair. In terms of linear algebra, this move replaces a vector space representing a single scale (the d-dimensional homology) with the image of a linear map between the vector spaces at two scales. To show that this actually works requires a combination of geometry, topology, algebra, and combinatorial representations. *A motivating goal when we started this work was to factor these concerns so as to make it easier to generalize.*

The main contributions of this paper are as follows.

1. A new proof of the TCC that applies to disconnected, compact domains with no smoothness assumptions satisfying only very weak topological conditions. In addition to being more general, our approach is simplified by factoring the geometric, topological, and algorithmic aspects of the original proof. We first give a completely geometric version of the main theorem that replaces all of the geometric assumptions on the domain with very weak topological assumptions. There is no need for the boundary of the domains to be smooth, or even manifold. This geometric TCC gives a guarantee directly about the coverage regions directly rather than previous work that required reasoning about the so-called shadow of the Rips complex.

2. An application of the TCC to weighted k-coverage via the weighted k-nearest neighbor distance. This introduces applications to a less restrictive class of distance-to-set functions permitted by the weakened assumption on the underlying space. This requires a slight modification to the original TCC algorithm where we replace the clique complexes with subcomplexes of their barycentric subdivisions.

1.1 Related Work Since de Silva & Ghrist developed the theory of coordinate-free coverage in homological sensor networks [17, 4, 5, 8], the theory has been extended in several directions. To test if coverage will be robust to failures, we have considered k-coverage. An alternative perspective on robustness was given by Munch et al. using probabilistic models of failure [14]. Dłotko et al. gave algorithms for computing the TCC in a distributed fashion [6]. Gamble et al. present some related methods for analyzing coverage in a dynamic setting, where sensors are allowed to move in time[9, 7]. Adams and Carlsson considered a related, dynamic version of the problem where instead of coverage, one wants to know if it is possible to evade a collection of moving

sensors [1].

2 Background

A **length space** is a metric space (M, d) in which the distance between any two points $x, y \in M$ is equal to the infimum of the lengths of all paths from x to y. Throughout we will consider a compact length space (M, d), where there is a homeomorphism from M to a compact subset of $\mathbb{R}^d$. To simplify notation we will interpret M directly as a subset of $\mathbb{R}^d$, but with the metric d on the points identified with M, rather than the Euclidean metric. We define the **complement** of a compact set $A \subseteq M$ as $\overline{A} := (\mathbb{R}^d \cup \{\infty\}) \setminus A$, where $\mathbb{R}^d \cup \{\infty\}$ is the compactification of $\mathbb{R}^d$, which we note is homeomorphic to the d-sphere.

If $A \subseteq M$ is endowed with weights $w_y \geq 0$ for all $y \in A$, the **weighted distance** from a point x to a weighted point y is defined as the power distance $\rho_y(x) := \sqrt{\mathrm{d}(x,y)^2 + w_y^2}$. Such a set is referred to as a **weighted set**. We use weighted distances to model coverage by disks of varying radii, where larger weights correspond to smaller radii. These weights can be used to model high dimensional noise when the weights of each point are defined to be the distance to its projection onto a domain.

Let $A \subseteq M$ be a weighted compact set. The **weighted k-nearest neighbor distance** from a point x to a weighted compact set $A \subseteq M$ is defined as

$$\mathrm{d}_k(x, A) := \inf_{K \in \binom{A}{k}} \max_{y \in K} \rho_y(x).$$

where $\binom{A}{k}$ denotes the collection of k-element subsets of A. The **weighted (k, ε)-offsets** of A are defined as

$$A_k^\varepsilon := \{x \in M \mid \mathrm{d}_k(x, A) \leq \varepsilon\}.$$

The **unweighted ε-offsets** of A are defined as

$$A^\varepsilon := \left\{x \in M \mid \min_{a \in A} \mathrm{d}(x, a) \leq \varepsilon\right\},$$

Note that for any weighted set A, any $\varepsilon \geq 0$, and any $k \geq 1$ we have $A_k^\varepsilon \subseteq A^\varepsilon$. The **coverage region** of a point $a \in A$ at scale ε is denoted $\mathrm{cov}(a, \varepsilon) = \{x \in M \mid \rho_a(x) \leq \varepsilon\}$.

Simplicial Complexes. A **simplicial complex** K is a collection of subsets, called **simplices**, of a vertex set V that is closed under taking subsets. That is, for all $\sigma \in K$ and $\tau \subset \sigma$ it must follow that $\tau \in K$. The **dimension** of a simplex $\sigma \in K$ is defined as $\dim(\sigma) := |\sigma| - 1$ where $|\cdot|$ denotes set cardinality. The dimension of a simplicial complex K is the maximum dimension of any simplex in K. We define a **pair of**

complexes to be a pair (K, L) where K is a simplicial complex and L is a subcomplex of K.

A **graph** $G = (V, E)$ is defined to be a simplicial complex of dimension at most 1, consisting of a vertex set V and an edge set $E \subseteq \binom{V}{2}$. Given a subset $U \subseteq V$ the **induced subgraph** in G by U is defined $G[U] := (U, E \cap \binom{U}{2})$.

Given a graph $G = (V, E)$ a **clique** is a collection of vertices $\sigma \subseteq V$ such that for all $u, v \in \sigma$ the edge $\{u, v\} \in E$. The **clique complex** of G is defined to be a simplicial complex with simplices for each clique in G.

$$\mathrm{Clq}(G) := \{\sigma \subseteq V \mid \forall u, v \in \sigma, \{u, v\} \in E\}$$

Given a pair of graphs (G, H) where H is a subgraph of G, we will denote the pair of Clique complexes as $\mathrm{Clq}(G, H) = (\mathrm{Clq}(G), \mathrm{Clq}(H))$.

The **Čech complex** of a finite collection of weighted points $A \subset M$ at scale ε is defined as

$$\check{C}^{\varepsilon}(A) := \left\{\sigma \subseteq A \mid \bigcap_{p \in \sigma} \mathrm{cov}(p, \varepsilon) \neq \emptyset\right\}.$$

The **(Vietoris-)Rips complex** of A at scale ε is defined as

$$\mathrm{Rips}^{\varepsilon}(A) := \left\{\sigma \subseteq A \mid \{p, q\} \in \check{C}^{\varepsilon}(A) \text{ for all } p, q \in \sigma\right\}.$$

The Rips complex is the clique complex derived from the edges in the Čech complex.

An important result about the relationship of Čech and Rips complexes follows from Jung's Theorem [11] relating the diameter of a point set A and the radius of the minimum enclosing ball:

$$(2.1) \qquad \check{C}^{\varepsilon}(A) \subseteq \mathrm{Rips}^{\varepsilon}(A) \subseteq \check{C}^{\vartheta_d \varepsilon}(A),$$

where the constant $\vartheta_d = \sqrt{\frac{2d}{d+1}}$ for unweighted sets and $\vartheta_d = 2$ for weighted sets (see [2]).

The k-Barycentric Decomposition. Given a simplicial complex K we define a **flag** in K to be an ordered subset of simplices $\{\sigma_1, \ldots, \sigma_t\} \subset K$ such that $\sigma_1 \subset \ldots \subset \sigma_t$. The **barycentric decomposition** of K is the simplicial complex formed by the set of flags of K and is defined as $\mathrm{Bary}(K) := \{U \subset K \mid U \text{ is a flag of } K\}$. The vertices of the barycentric decomposition are the simplices of K. We define the **degree** of a flag $\sigma_1 \subset \cdots \subset \sigma_t$ to be $|\sigma_1|$. The k-**barycentric decomposition** of a complex S is defined as

$$k\text{-}\mathrm{Bary}(K) := \{U \subset K \mid U \text{ is a flag in } K \text{ with } |U| \geq k\}.$$

The k-barycentric decomposition of the Clique complex of a graph G will be denoted

$$\mathrm{Clq}_k(G) = k\text{-}\mathrm{Bary}(\mathrm{Clq}(G)).$$

Similarly, the k-barycentric decomposition of the Čech complex of a finite point set A at a scale ε will be denoted

$$\check{C}_k^{\varepsilon}(A) = k\text{-}\mathrm{Bary}(\check{C}^{\epsilon}(A)).$$

$$K \qquad\qquad 0\text{-}\mathrm{Bary}(K)$$

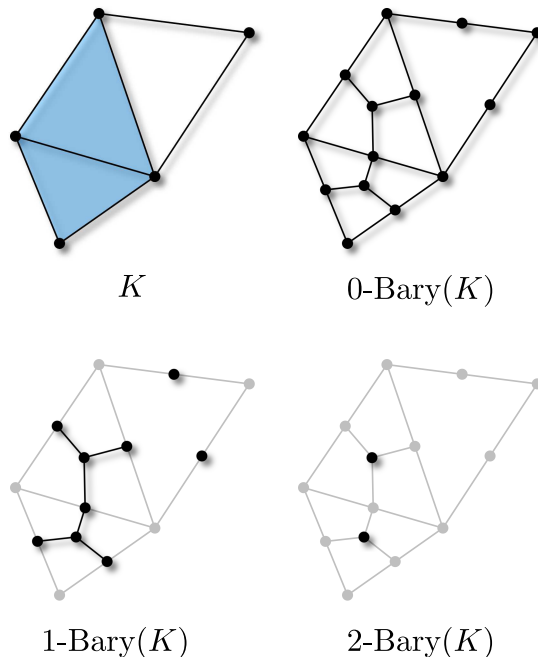$$1\text{-}\mathrm{Bary}(K) \qquad\qquad 2\text{-}\mathrm{Bary}(K)$$

Figure 1: A simplicial complex K and its $0, 1,$ and 2-barycentric decompositions.

Homology and Persistent Homology. Homology is a tool from algebraic topology that gives a computable signature for a shape that is invariant under many topological equivalences, in particular homemorphisms and homotopy equivalences. It gives a way to quantify the components, loops, and voids in a topological space. It is a favored tool for applications because its computation can be phrased as a matrix reduction problem with matrices representing a finite simplicial complex.

Throughout, we assume singular homology over a field, so the kth homology group $\mathrm{H}_k(C)$ of a space C is vector space. When considering the homology groups of all dimensions, we will write $\mathrm{H}_*(C)$. We will make extensive use of relative homology. That is, for a pair of spaces (A, B) with $B \subseteq A$, we write $\mathrm{H}_*(A, B)$ for the homology of A relative to B.

We can also talk about the homology of a map between two spaces. Given two spaces A and B and a map $f : A \rightarrow B$, we can consider the homology of both the spaces and the map for all homology groups due to the functoriality of homology, i.e. we have a map $f_* = H_*(A) \rightarrow H_*(B)$, which we will denote

$f_* := H_*(A \to B)$. Of particular interest in this work is the homology map induced by inclusion from one space to another, in which commutativity of diagrams of spaces is preserved when passed to a diagram of homology groups.

Other Topological Notions. We will employ several other standard notions from topology. A space is **triangulable** if it is homeomorphic to a finite simplicial complex. Triangulability acts as a non-degeneracy condition.

There are vector spaces dual to the homology groups called the **cohomology groups** and they are denoted with superscripts as $H^*(C)$ with respect to a space C. For finite-dimensional homology groups, the so-called **Universal Coefficient Theorem** implies that the r-dimensional homology and cohomology groups are isomorphic. This will allow us to switch between the two theories when it is convenient.

The other way we will switch between homology and cohomology is by **Alexander duality** which states, in general, that for pairs of nonempty compact spaces in $\mathbb{R}^d \cup \{\infty\}$, their r-dimensional relative homology is isomorphic to their complement spaces' $(d-r)$-dimensional relative cohomology, i.e. $H_r(X, Y) \cong H^{d-r}(\overline{Y}, \overline{X})$,. The specific version we use will be discussed in Section 5.

Nerves and Persistent Nerves. The Čech complex is a special case of a general construction known as a **nerve**. Let $U = \{U_i \mid i \in I\}$ be a collection of sets, where I is any indexing set. The nerve of U is the simplicial complex with vertex set I such that $\sigma \subseteq I$ is a simplex if and only if $\bigcap_{i \in \sigma} U_i \neq \emptyset$. We say that U **covers** the set $\bigcup_{i \in I} U_i$ and it is a **good cover** if the intersections are empty or contractible. For such covers, one can relate the nerve of the cover and union using the so-called Nerve Theorem. Chazal and Oudot generalized the nerve theorem to the persistence setting [3] and Sheehy extended it to k-coverage [16]. This is captured in the following lemma, where $(\check{C}_k^\varepsilon(A, B)) = (\check{C}_k^\varepsilon(A)), (\check{C}_k^\varepsilon(B))$.

LEMMA 2.1. *For any $B \subset A \subseteq M$, if the coverage regions $\{\operatorname{cov}(a, \alpha) \mid a \in A\}$ form a good cover of A_1^α and similarly for A_1^β, B_1^α, and B_1^β, then the following diagram commutes for all k and the vertical maps are isomorphisms.*

$$
\begin{array}{ccc}
H_*(A_k^\alpha, B_k^\alpha) & \longrightarrow & H_*(A_k^\beta, B_k^\beta) \\
\Big\downarrow{\cong} & & \Big\downarrow{\cong} \\
H_*(\check{C}_k^\alpha(A, B)) & \longrightarrow & H_*(\check{C}_k^\beta(A, B))
\end{array}
$$

A combinatorial construction of this fact appears in [16], but a more direct topological argument can be found in Appendix B.

3 Assumptions

Strange examples abound in topology. One must make some assumptions about the underlying domain to make the TCC possible. In this section, we will first introduce and illustrate the minimal geometric properties that we require of the bounded domain to be covered. We will weaken the geometric and input assumptions on the domain and point sample from those required for the topological coverage criterion of de Silva & Ghrist to apply it a much wider class of sets.

For a pair of sets $(\mathcal{D}, \mathcal{B})$ such that $\mathcal{B} \subseteq \mathcal{D}$, we say that $\mathcal{B}$ *surrounds* $\mathcal{D}$ if there is no path from $\mathcal{D} \setminus \mathcal{B}$ to $\overline{\mathcal{D}}$ that does not intersect $\mathcal{B}$. Formally, $\mathcal{B}$ surrounds $\mathcal{D}$ if and only if $H_0(\mathcal{D} \setminus \mathcal{B}) \cong H_0(\overline{\mathcal{B}}, \overline{\mathcal{D}})$. If such a pair satisfies the following conditions for $0 < 3\alpha \leq \beta$, we want to certify that a finite sample $P \subset \mathcal{D}$ covers $\mathcal{D}$ at scale α in the sense that $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subset P^\alpha$.

Geometric Assumptions

0. *(The Domain)* $\mathcal{D}$ is a bounded, compact length space homeomorphic to a subset of $\mathbb{R}^d$ and $\mathcal{B} \subseteq \mathcal{D}$ is closed and surrounds $\mathcal{D}$.

1. *(Components are not too small)* The map $H_0(\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta} \hookrightarrow \mathcal{D} \setminus \mathcal{B}^{2\alpha})$ is surjective.

2. *(Components are not too close)* The map $H_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow (\mathcal{D} \setminus \mathcal{B})^{2\alpha})$ is injective.

Lemma 3.1 allows us talk about the homology of a subset of the domain $\mathcal{D} \setminus \mathcal{B}^\varepsilon$ in terms of relative homology.

LEMMA 3.1. *If $\mathcal{B}$ surrounds $\mathcal{D}$, then for all $\varepsilon > 0$, $H_0((\mathcal{D} \setminus \mathcal{B}^\varepsilon, \emptyset) \hookrightarrow (\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon}))$ is an isomorphism.*

Proof. First we show the map is injective. Given some non-trivial 0-chain $[x] \in H_0(\mathcal{D} \setminus \mathcal{B}^\varepsilon)$, we can pick a representative point $x \in \mathcal{D} \setminus \mathcal{B}^\varepsilon \subseteq \overline{\mathcal{B}^\varepsilon}$. Because $\mathcal{B}$ surrounds $\mathcal{D}$, there does not exist a path from $\mathcal{D} \setminus \mathcal{B}$ to $\overline{\mathcal{D}}$ that does not intersect $\mathcal{B}$, and so there does not exist a path from $\mathcal{D} \setminus \mathcal{B}^\varepsilon$ to $\overline{\mathcal{D}^\varepsilon}$ that does not intersect $\mathcal{B}^\varepsilon$. Thus, $[x] \neq 0 \in H_0(\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon})$.

Next we show the map is surjective. Any $[x] \in H_0(\overline{\mathcal{B}^\varepsilon}, \overline{\mathcal{D}^\varepsilon})$, is represented by a point x in a connected component of $\overline{\mathcal{B}^\varepsilon} \setminus \overline{\mathcal{D}^\varepsilon} = \mathcal{D} \setminus \mathcal{B}^\varepsilon$, and thus a homology class $[x] \in H_0(\mathcal{D} \setminus \mathcal{B}^\varepsilon)$.

Assumption 1 disallows domains with components that are too small to be included in the map from $\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta} \hookrightarrow \mathcal{D} \setminus \mathcal{B}^{2\alpha}$ so we can reliably compare the coverage region to the sampled subset of the domain in terms of the 0-dimensional homology, or connected components. This approach is motivated by the notion
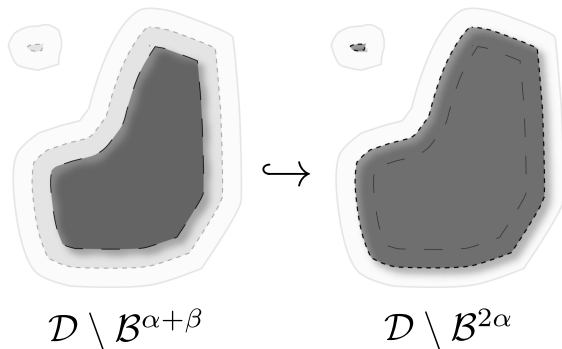
that a coverage hole in a d-dimensional domain can be represented by a connected component of the complement space under Alexander duality. Fig. 2 illustrates a domain in which the induced map is not surjective.

Assumption 2 requires that the components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ are spaced out well enough so that no components are joined with inclusion into $\mathcal{D}^{2\alpha}$. Fig. 3 illustrates domains which violate Assumption 2. Assumptions 0-2 are necessary in order to prove the Geometric TCC, and as we will see Assumption 2 is used to bound the number of connected components of the shrunken domain in the Algorithmic TCC.

Let $(\mathcal{D}, \mathcal{B})$ be a pair of spaces satisfying Assumptions 0–2 for constants $\alpha > 0$, $\beta \geq 3\alpha$. The input to Algorithm 1 will be a pair of graphs (G_1, G_2), a finite weighted point sample $P \subset \mathcal{D}$ and a subsample $Q = \{p \in P \mid \mathrm{cov}(p, \alpha) \cap \mathcal{B} \neq \emptyset\}$.

Figure 2: A domain that violates Assumption 1 as the upper-left component appears in the inclusion from $\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta}$ to $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$.
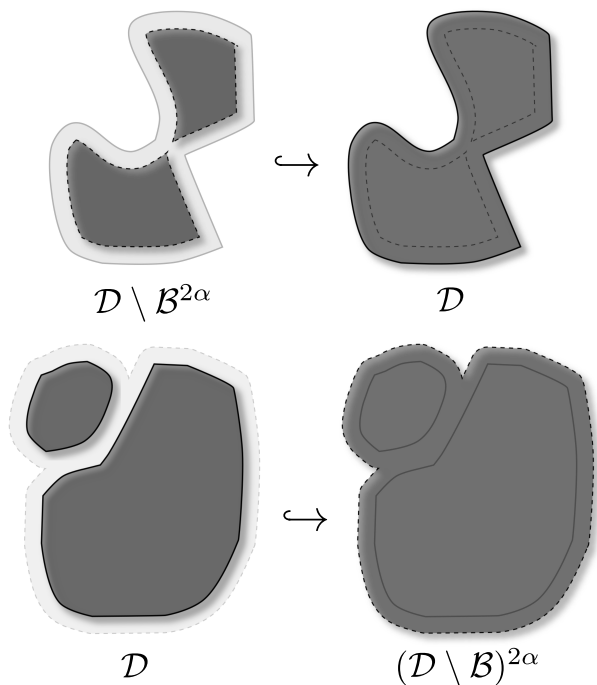

Figure 3: Domains which violate Assumption 2 in which components are lost in the inclusions $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow \mathcal{D}$ and $\mathcal{D} \hookrightarrow (\mathcal{D} \setminus \mathcal{B})^{2\alpha}$ respectively. In the first case a single component is pinched apart in $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. In the second, two components which are too close are merged in $(\mathcal{D} \setminus \mathcal{B})^{2\alpha}$.

Input Assumptions

3. The graphs G_1, G_2 have a vertex set $P \subset \mathcal{D}$ and subgraphs $G_1[Q], G_2[Q]$ induced by restriction to the vertex set $Q = \{p \in P \mid \mathrm{cov}(p, \alpha) \cap \mathcal{B} \neq \emptyset\}$.

4. $\mathcal{U} = \{\mathrm{cov}(p, \varepsilon) \mid p \in P\}$ is a good cover for $\varepsilon \in \{\alpha, \beta\}$.

5. $\mathrm{Clq}_k(G_1) \subseteq \check{C}_k^\alpha(P) \subseteq \check{C}_k^\beta(P) \subseteq \mathrm{Clq}_k(G_2)$.

6. Each component of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ contains a point in P.

7. There exists a triangulation $\mathcal{K}$ of $\mathbb{R}^d \cup \{\infty\}$ and triangulations of P_k^ε and Q_k^ε, L_ε and M_ε respectively, where $M_\varepsilon \subset L_\varepsilon$ in $\mathcal{K}$, for $\varepsilon \in \{\alpha, \beta\}$.

The input graphs will be used to construct clique complexes, which by Assumption 5 can be interleaved with Čech complexes at scale α and β, and then Assumption 4 allows us to apply the Persistent Nerve Lemma [3]. Note that Assumption 5 is satisfied when every clique $\sigma \subseteq G_1$ is such that $\bigcap_{v \in \sigma} \mathrm{cov}(v, \alpha) \neq \emptyset$ and $\mathrm{cov}(u, \beta) \cap \mathrm{cov}(v, \beta) \neq \emptyset$ implies $\{u, v\} \in E(G_2)$ for all $u, v \in \sigma$. When the domain $\mathcal{D}$ is equipped with the Euclidean metric, Assumption 5 is equivalent to the interleaving provided by Jung's theorem (Equation 2.1) where the k-clique complex of G_1 can be taken as the k-Rips complex of P at scale α/ϑ_d. Finally, Assumption 6 is in the proof of Lemma 5.2 to bound the number of connected components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ using a computable combinatorial structure.

Relationship to the de Silva & Ghrist TCC According to de Silva and Ghrist [4, Remark 4.5], the smooth manifold hypothesis seems to be a necessary requirement in order to apply the TCC. Because their analysis involved directly comparing the thickened re-
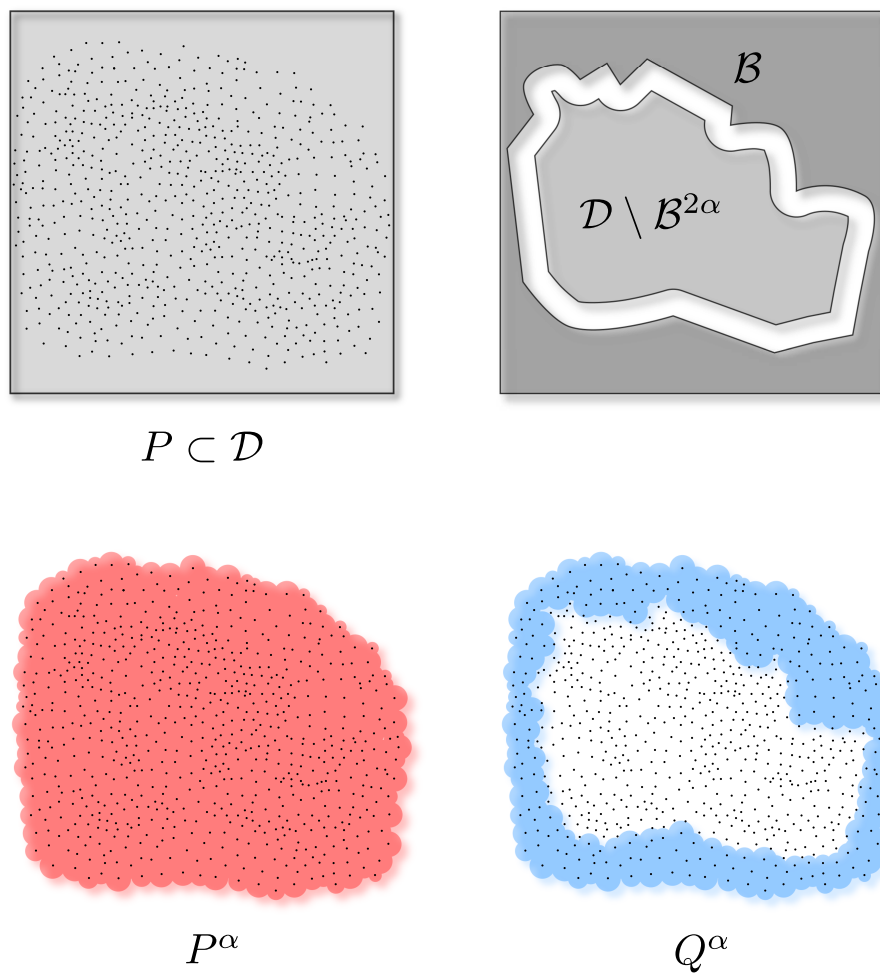
Figure 4: An example of a pair $(\mathcal{D}, \mathcal{B})$ in which $\mathcal{B}$ surrounds $\mathcal{D}$. We would like to assert conditions which allow us to verify a weighted sample P k-covers $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ at scale α. We can then compare the relative homology of the pair (P^α, Q^α) to $(\mathcal{D}, \mathcal{B}^{2\alpha})$.

gion around the boundary to the boundary points' corresponding complex, it was necessary for them to show that the thickness of this region is such that any topological noise in the complex is eliminated with inclusion from scale α to scale β. This amounts to proving that that cycles lying entirely in a thickening of the boundary of the domain cannot persist, as shown in [4, Lemma 3.3]. In Fig. 5, we show a domain without a smooth boundary in which the thickened boundary contains a cycle that persists across a range of scales. This example illustrates the the contrast between [4, Lemma 3.3] and our Lemma 4.1, which requires only Assumptions 0 and 1, and in which the persistence of this cycle does not indicate the persistence of a relative cycle. This is the critical idea to understand how we show the correctness of the TCC under these weaker assumptions.

4   A New Proof of the Topological Coverage Criterion

Consider a domain $\mathcal{D}$, a set $\mathcal{B}$ that surrounds $\mathcal{D}$, and constants α and β such that $0 < 3\alpha \leq \beta$ satisfying Assumptions 0–2. Let $P \subset \mathcal{D}$ be a weighted finite point set and $Q = \{p \in P \mid \mathrm{cov}(p, \alpha) \cap \mathcal{B} \neq \emptyset\}$ be a subsample of P within distance α of the boundary. We will give a sufficient condition to guarantee coverage of a shrunken domain $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ by P_k^{α}, i.e. k-coverage.

We will assume non-negative weights $w_x \geq 0$ assigned to each $x \in P$, and that $w_x = 0$ for all points $x \in \mathcal{D} \setminus P$. This implies that $\mathcal{D}_k^{\varepsilon} = \mathcal{D}^{\varepsilon}$, and similarly $\mathcal{B}_k^{\varepsilon} = \mathcal{B}^{\varepsilon}$, so we will simply use the notation $\mathcal{D}^{\varepsilon}$ and $\mathcal{B}^{\varepsilon}$ throughout. Moreover, we know that $P_k^{\alpha} \subseteq \mathcal{D}_k^{\alpha} = \mathcal{D}^{\alpha}$ by the monotonicity of d_k. For any arbitrary weighted compact set $A \subseteq \mathcal{D}$, $A_k^{\varepsilon} \subseteq A_1^{\varepsilon} \subseteq A^{\varepsilon}$ and $Q \subseteq \mathcal{B}^{\alpha}$, so for $\varepsilon \geq 0$, $Q_k^{\varepsilon} \subseteq Q^{\varepsilon} \subseteq \mathcal{B}^{\alpha+\varepsilon}$. Diagram (4.2) is a commutative diagram of inclusion maps. Complementing all the spaces and reversing inclusions give Diagram (4.3).

$$(4.2) \qquad \begin{array}{ccc} (P_k^{\alpha}, Q_k^{\alpha}) & \hookrightarrow & (P_k^{\beta}, Q_k^{\beta}) \\ \downarrow & & \downarrow \\ (\mathcal{D}^{2\alpha}, \mathcal{B}^{2\alpha}) & \hookrightarrow & (\mathcal{D}^{\alpha+\beta}, \mathcal{B}^{\alpha+\beta}) \end{array}$$

$$(4.3) \qquad \begin{array}{ccc} (\overline{\mathcal{B}^{\alpha+\beta}}, \overline{\mathcal{D}^{\alpha+\beta}}) & \overset{j}{\hookrightarrow} & (\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) \\ \uparrow & & \uparrow \\ (\overline{Q_k^{\beta}}, \overline{P_k^{\beta}}) & \overset{i}{\hookrightarrow} & (\overline{Q_k^{\alpha}}, \overline{P_k^{\alpha}}) \end{array}$$

We use the diagram of the complements so that all the proofs only require arguments about connectivity

(0-dimensional homology). Diagram (4.4) is formed from Diagram (4.3) by applying the homology functor.

$$(4.4) \qquad \begin{array}{ccc} \mathrm{H}_0(\overline{\mathcal{B}^{\alpha+\beta}}, \overline{\mathcal{D}^{\alpha+\beta}}) & \overset{j_*}{\longrightarrow} & \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) \\ \downarrow & & \downarrow \\ \mathrm{H}_0(\overline{Q_k^{\beta}}, \overline{P_k^{\beta}}) & \overset{i_*}{\longrightarrow} & \mathrm{H}_0(\overline{Q_k^{\alpha}}, \overline{P_k^{\alpha}}) \end{array}$$

This commutative diagram induces the most important map in this paper, $p_* : \mathrm{im}\, j_* \to \mathrm{im}\, i_*$. Though reversed and complemented, this map describes the topology of the offsets embedded into the domain, where the scale change eliminates noise. It will capture exactly the topological information we want. Analyzing p_* directly simplifies our approach compared to previous work, simplifying the proof and eliminating some hypotheses.

Theorem 4.1 proves that we can infer coverage from the rank of i_* in order to construct an algorithm confirming coverage by comparing the rank of i_* to the number of connected components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. This motivates Assumption 1 ensuring that $\mathrm{im}\, j_*$ reflects the 0-dimensional homology of our coverage domain $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$.

The following two lemmas prove two important properties of p_*. First, we show that Assumptions 0 and 1 suffice to guarantee that p_* is surjective (Lemma 4.1). Then, we prove that if p_* is injective, then the domain is covered (Lemma 4.2). These will lead directly to the geometric form of the Topological Coverage Criterion (Theorem 4.1).

LEMMA 4.1. *Given Assumptions* 0 *and* 1*, the map* p_* *is surjective.*

Proof. Assumption 1 and Lemma 3.1 imply that j_* is surjective. First we choose a basis for $\mathrm{im}\, i_*$ such that each basis element is represnted by a point in $P_k^{\alpha} \setminus Q_k^{\beta}$. Consider $x \in P_k^{\alpha} \setminus Q_k^{\beta}$ such that $[x] \neq 0 \in \mathrm{im}\, i_*$.

Suppose $x \in \mathcal{B}^{2\alpha}$. By definition of the offsets, there is a point $y \in \mathcal{B}$ such that $\mathrm{d}(x, y) \leq 2\alpha$. Because $x \in \overline{Q_k^{\beta}}$ by hypothesis, $\mathrm{d}_k(x, Q) > \beta$. We will show that if a point z is in the shortest path $\overline{xy}$, then $z \in \overline{Q_k^{\alpha}}$. For any $z \in \overline{xy}$, we have $\mathrm{d}(x, z) \leq \mathrm{d}(x, y) \leq 2\alpha$, thus the following inequality holds.

$$\begin{aligned} \mathrm{d}_k(z, Q) &\geq \mathrm{d}_k(x, Q) - \mathrm{d}(x, z) && [\mathrm{d}_k \text{ is Lipschitz}] \\ &> \beta - 2\alpha && [\mathrm{d}_k(x, Q) > \beta, \mathrm{d}(x, z) \leq 2\alpha] \\ &\geq \alpha. && [\beta \geq 3\alpha] \end{aligned}$$

From this inequality we conclude that $z \in \overline{Q_k^{\alpha}}$ for all $z \in \overline{xy}$, and thus $\overline{xy} \subseteq \overline{Q_k^{\alpha}}$ and $y \in \mathcal{B} \cap \overline{Q_k^{\alpha}}$. Now suppose $y \in P_k^{\alpha}$ as well. By the definition of P_k^{α}, for some $A \in \binom{P}{k}$, $y \in \mathrm{cov}(p, \alpha)$ for all $p \in A$, which implies
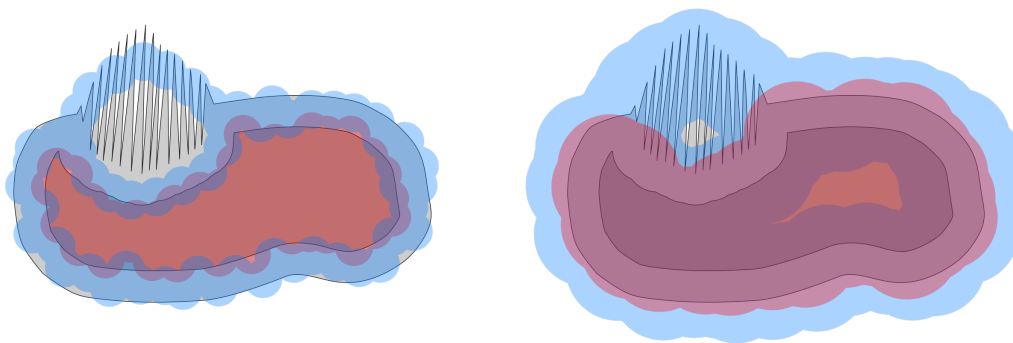
Figure 5: This instance illustrates the failure of Lemma 3.3 of [4] when the boundary is not smooth. A cycle that is trivial in the thickened boundary persists. This highlights the need to work with the relative homology of the domain modulo the boundary rather than the homology of the boundary alone. Such a cycle in the boundary cannot form a relative cycle.

that $A \subseteq Q$, so $y \in Q_k^\alpha$, a contradiction. Thus we may conclude that $y \notin P_k^\alpha$, which is equivalent to $y \in \overline{P_k^\alpha}$.

Any path $\gamma : [0,1] \to \overline{Q_k^\alpha}$ such that $\gamma(0) = x$ and $\gamma(1) = y$, generates a class $[\gamma]$ in the chain group $C_1(\overline{Q_k^\alpha})$ containing γ. For $[\gamma] \in C_1(\overline{Q^\alpha}_k, \overline{P_k^\alpha})$ it follows $\partial([\gamma]) = [x + y] = [x]$ as $y \in \overline{P_k^\alpha}$, and therefore that there must exist $z \in \overline{xy} \cap Q_k^\alpha$. This is a contradiction as we have shown that $\overline{xy} \cap Q_k^\alpha = \emptyset$, and thus x cannot be in $\mathcal{B}^{2\alpha}$.

Now, we may assume $x \in \overline{\mathcal{B}^{2\alpha}}$. Then $x \in \mathcal{D} \setminus \mathcal{B}^{2\alpha}$ so $[x] \neq 0 \in H_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}})$. Because j_* is surjective, $H_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) = \operatorname{im} j_*$ and thus $p_*([x]) = [x]$ and so $[x] \in \operatorname{im} p_*$. It follows that p_* is surjective. $\square$

The following lemma will allow us to confirm coverage by comparing the ranks of $\operatorname{im} i_*$ and $\operatorname{im} j_*$.

LEMMA 4.2. *Given Assumptions 0 and 1, if p_* is injective then $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$.*

Proof. The proof is essentially the same as that presented by de Silva & Ghrist [4]. We include it here in our own notation for completeness.

Suppose for contradiction p_* is injective and there exists a point $x \in (\mathcal{D} \setminus \mathcal{B}^{2\alpha}) \setminus P_k^\alpha$. Thus $[x] \neq 0 \in H_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}})$, because the point x is in some connected component of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. Moreover, $[x] \in \operatorname{im} j_*$, because j_* is surjective by Assumption 1. Consider the following sequence of maps induced by inclusions.

$$H_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) \xrightarrow{f_*} H_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}} \cup \{x\}) \xrightarrow{g_*} H_0(\overline{Q_k^\alpha}, \overline{P_k^\alpha})$$

As $f_*([x])$ is zero in $H_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}} \cup \{x\})$, then $p_*([x]) = (g_* \circ f_*)([x]) = 0$, contradicting the assumption that p_* is injective. $\square$

We are now ready to state and prove the geometric version of the TCC, relating the k-covered offsets of the sample to the underlying domain. Note that it will not directly give an algorithm (that will come in Theorem 5.2), but instead a result about the offsets. This differs from previous work that analyzed an embedding of a Rips complex.

THEOREM 4.1. (THE GEOMETRIC TCC) *Let $(\mathcal{D}, \mathcal{B})$ be a pair of sets satisfying Assumptions 0 and 1. Let α and β be constants such that $0 < 3\alpha \leq \beta$. Let $P \subset \mathcal{D}$ be a finite set with $Q = \{p \in P \mid \operatorname{cov}(p, \alpha) \cap \mathcal{B} \neq \emptyset\}$. Let i_* and j_* be the maps in Diagram (4.4). If $\operatorname{rk} i_* \geq \operatorname{rk} j_*$ then $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$.*

Proof. Lemma 4.1 implies that $p_* : \operatorname{im} j_* \to \operatorname{im} i_*$ is surjective, so $\operatorname{rk} i_* \leq \operatorname{rk} j_*$. So with the assumption that $\operatorname{rk} i_* \geq \operatorname{rk} j_*$, $\operatorname{rk} i_* = \operatorname{rk} j_*$. Since P is a finite point set, $\operatorname{im} i_*$ is finite-dimensional and by equality, $\operatorname{im} j_*$ is as well, so p_* is an isomorphism, and thus it is injective. Lemma 4.2 then implies $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$. $\square$

5 Computing the TCC

In the previous section we proved sufficient conditions for generalized coverage in terms of the offsets of the input points. In this section, we provide an algorithm to check for k-coverage of a shrunken domain by a weighted sample P, i.e. that $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$ and a proof of its correctness.

Given a pair $(\mathcal{D}, \mathcal{B})$ and non-negative constants α, β such that $\beta \geq 3\alpha$ satisfying Assumptions 0–2, the input to Algorithm 1 is: a pair of graphs (G_1, G_2), a finite weighted point sample $P \subset \mathcal{D}$, a set $Q = \{p \in P \mid \operatorname{cov}(p, \alpha) \cap \mathcal{B} \neq \emptyset\}$, and $k > 0$ satisfying Assumptions 3–7. G_1 and G_2 ideally represent the

connectivity graphs of the points in P at two scales, and will be used to construct clique complexes that can be related to the Geometric TCC within a constant factor. Note that if the algorithm's output is false it does not necessarily mean there is not coverage.

Algorithm 1 Check if $\mathcal{D} \setminus B^{2\alpha} \subseteq P_k^\alpha$

1: **procedure** K-COVERAGE(G_1, G_2, P, Q, k)
2:    let $c := |Components(G_1[P \setminus Q])|$
3:    let

$$r := \mathrm{rk}\ \mathrm{H}_d(\mathrm{Clq}_k(G_1, G_1[Q]) \hookrightarrow \mathrm{Clq}_k(G_2, G_2[Q]))$$

4:    **if** $c = r$ **then return** True
5:    **else return** False

Lemma 5.1, Lemma 5.2 and Theorem 5.2 collectively provide a proof of correctness of Algorithm 1. Lemma 5.1 bounds the rank of the map between the clique complexes at different scales by rk i_*, in order to compare it to rank j_* through Theorem 4.1. Lemma 5.2 states that if the components are separated enough and each component of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$ contains a point in P, formally defined in Assumptions 2 and 6, then the number of connected components of the graph induced by restriction to $P \setminus Q$ on G_1 provides an upper bound for the number of components of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$.

In order to prove Lemma 5.1, we must invoke Alexander Duality. The form which is most useful based on our triangulability assumptions, Assumption 7, is the following. The naturality of it implies that homology maps induced by inclusion maps at the level of pairs of spaces commute with the duality.

THEOREM 5.1. (ALEXANDER DUALITY [10]) *Let* K *be an abstract simplicial complex which is a combinatorial oriented d-manifold. Given a pair (L, M), where L is a subcomplex of some refinement of K, and M is a subcomplex of L, there is a natural isomorphism*

$$\mathrm{H}^r(L, M) \to \mathrm{H}_{d-r}(\overline{M}, \overline{L}).$$

The complements of the simplicial complexes in Theorem 5.1 are the simplicial complexes in K that do not share vertices with the original complex. By considering a sufficiently refined triangulation $\mathcal{K}$ of $\mathbb{R}^d \cup \{\infty\}$ and the corresponding triangulations of P_k^α, Q_k^α, P_k^β, and Q_k^β, whose existence is assumed in Assumption 7, Theorem 6.2.17 in Spanier [18], Theorem 74.1 in Munkres [15] and Theorem 5.2 in [10] collectively provide the following corollary to Theorem 5.1.

COROLLARY 5.1. *Given $\alpha, \beta > 0$, and P and Q satisfying Assumption 7, we have the following two natural*

isomorphisms.

$$\mathrm{H}_0(\overline{Q_k^\beta}, \overline{P_k^\beta}) \to \mathrm{H}^d(P_k^\beta, Q_k^\beta)$$

and

$$\mathrm{H}_0(\overline{Q_k^\alpha}, \overline{P_k^\alpha}) \to \mathrm{H}^d(P_k^\alpha, Q_k^\alpha).$$

LEMMA 5.1. *The rank of the map* $\mathrm{H}_d(\mathrm{Clq}_k(G_1, G_1[Q])$ $\hookrightarrow$ $\mathrm{Clq}_k(G_2, G_2[Q]))$ *induced by inclusion is at most* rk i_*.

Proof. Corollary 5.1 with respect to Diagram (4.4) implies that $\mathrm{rk}(\mathrm{H}^d(P_k^\beta, Q_k^\beta) \to \mathrm{H}^d(P_k^\alpha, Q_k^\alpha)) = \mathrm{rk}\ i_*$. By the Universal Coefficient Theorem, we then have $\mathrm{rk}(\mathrm{H}_d(P_k^\alpha, Q_k^\alpha) \to \mathrm{H}^d(P_k^\beta, Q_k^\beta)) = \mathrm{rk}\ i_*$, and by Lemma 2.1, for all $\varepsilon \geq 0$, $\mathrm{H}_*(P_k^\varepsilon, Q_k^\varepsilon) \cong \mathrm{H}_*(\check{C}_k^\varepsilon(P, Q))$, so $\mathrm{rk}(\mathrm{H}_d(\check{C}_k^\alpha(P, Q) \hookrightarrow \check{C}_k^\beta(P, Q))) = \mathrm{rk}\ i_*$.

By Assumption 5, the inclusion $\mathrm{Clq}_k(G_1, G_1[Q]) \hookrightarrow \mathrm{Clq}_k(G_2, G_2[Q])$ can be factored as

$$\mathrm{Clq}_k(G_1, G_1[Q]) \hookrightarrow \check{C}_k^\alpha(P, Q) \hookrightarrow \check{C}_k^\beta(P, Q) \hookrightarrow \mathrm{Clq}_k(G_2, G_2[Q]).$$

It follows that

$$\mathrm{rk}(\mathrm{H}_d(\mathrm{Clq}_k(G_1, G_1[Q]) \hookrightarrow \mathrm{Clq}_k(G_2, G_2[Q])))$$
$$\leq \mathrm{rk}(\mathrm{H}_d(\check{C}_k^\alpha(P, Q) \hookrightarrow \check{C}_k^\beta(P, Q)))$$
$$= \mathrm{rk}\ i_*.$$

LEMMA 5.2. *Let $(\mathcal{D}, \mathcal{B})$ be a pair of spaces satisfying Assumptions 0 and 2 for constants $\alpha, \beta \geq 3\alpha > 0$. If $P \subset \mathcal{D}$, $Q = \{p \in P \mid \mathrm{cov}(p, \alpha) \cap \mathcal{B} \neq \emptyset\}$ and the graph G_1 satisfy Assumptions 3, 5 and 6 for some constant k then $|Components(G_1[P \setminus Q])| \geq \mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha})$.*

Proof. Assume there exists $p, q \in P \setminus Q$ such that p and q are connected in $\mathrm{Clq}(G_1[P \setminus Q])$, but not in $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. By Assumption 5, we have that $\mathrm{d}(p, q) \leq 2\alpha$ and $[p] \neq [q]$ in $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha})$. However, the shortest path $\overline{pq} \in (\mathcal{D} \setminus \mathcal{B})^{2\alpha}$, as the distance between p and q is less than 2α, so $[p] = [q]$ in $\mathrm{H}_0(\mathcal{D}^{2\alpha})$, which implies that $\mathrm{H}_0(\mathcal{D} \setminus \mathcal{B}^{2\alpha} \hookrightarrow (\mathcal{D} \setminus \mathcal{B})^{2\alpha})$ is not injective, a contradiction to Assumption 2.

THEOREM 5.2. (ALGORITHMIC TCC) *Consider a pair $(\mathcal{D}, \mathcal{B})$, a finite point sample $P \subset \mathcal{D}$, and constants k, α, β, where $0 < 3\alpha \leq \beta$, satisfying Assumptions 0–7. If*

$$\mathrm{rk}\ \mathrm{H}_d(\mathrm{Clq}_k(G_1, G_1[Q]) \hookrightarrow \mathrm{Clq}_k(G_2, G_2[Q]))$$
$$= |Components(G_1[P \setminus Q])|$$

then $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$.

Proof. For simplicity, define

$$a_* := \mathrm{H}_d(\mathrm{Clq}_k(G_1, G_1[Q]) \hookrightarrow \mathrm{Clq}_k(G_2, G_2[Q]))$$

and set $c = |\text{Components}(G_1[P \setminus Q])|$, $m = \text{H}_0(D \setminus \mathcal{B}^{2\alpha})$. By our hypotheses and Lemma 5.1, rk $i_* \geq$ rk $a_* = c$. By Lemma 5.2, $c \geq m$, and Assumption 2 implies that j_* is surjective by Lemma 3.1 so by definition of $\mathcal{B}$ surrounding $\mathcal{D}$, $m = $ rk j_*. Thus rk $i_* \geq$ rk $a_* = c \geq m = $ rk j_*, namely rk $i_* \geq$ rk j_*, so by Theorem 4.1 we can conclude $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P_k^\alpha$.

References

[1] Henry Adams and Gunnar Carlsson. Evasion paths in mobile sensor networks. *International Journal of Robotics Research 34 (2014)*, pages 90–104, 2013.

[2] Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 168–180, 2015.

[3] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SCG '08, pages 232–241, New York, NY, USA, 2008. ACM.

[4] Vin de Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7:339–358, 2007.

[5] Vin de Silva and Robert Ghrist. Homological sensor networks. *Notices Amer. Math. Soc.*, 54(1):10–17, 2007.

[6] Paweł Dłotko, Robert Ghrist, Mateusz Juda, and Marian Mrózek. Distributed computation of coverage in sensor networks by homological methods. *Applicable Algebra in Engineering, Communication and Computing*, 23:29–58, 2012.

[7] Jennifer Gamble, Harish Chintakunta, and Hamid Krim. Coordinate-free quantification of coverage in dynamic sensor networks. *CoRR*, abs/1411.7337, 2014.

[8] Robert Ghrist and Abubakr Muhammad. Coverage and hole-detection in sensor networks via homology. In *Fourth International Conference on Information Processing in Sensor Networks (IPSN'05), UCLA*, pages 254–260, 2005.

[9] Harish Chintakunta Jennifer Gamble and Hamid Krim. Applied topology in static and dynamic sensor networks. *Signal Processing and Communications (SP-COM), 2012 International Conference*, pages 1–5, 2011.

[10] William Julian, Ray Mines, and Fred Richman. Alexander duality. *Pacific J. Math*, 106(1):115–127, 1983.

[11] Heinrich Jung. Über die kleinste kugel, die eine räumliche figur einschließt. *J. Reine Angew. Math*, 123:214–257, 1901.

[12] Dmitry Koslov. *Combinatorial Algebraic Topology*. Springer, 2000.

[13] Ryan H. Lewis and Dmitriy Morozov. Parallel computation of persistent homology using the blowup complex. In *Proceedings of the Annual Symposium on Parallelism in Algorithms and Architectures*, pages 323–331, 2015.

[14] Elizabeth Munch, Michael Shapiro, and John Harer. Failure filtrations for fenced sensor networks. *International Journal of Robotics Research*, 2012.

[15] James R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley Publishing Company, 1984.

[16] Donald R. Sheehy. A multicover nerve for geometric inference. In *The Canadian Conference in Computational Geometry*, pages 309–314, 2012.

[17] Vin De Silva and Robert Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *International Journal of Robotics Research*, 25:1205–1222, 2006.

[18] Edwin H. Spanier. *Algebraic Topology*. Springer New York, 1966.

[19] Afra Zomorodian and Gunnar Carlsson. Localized homology. *Computational Geometry: Theory and Applications*, 41(3):126–148, 2008.

A   Distance Lemmas

Our use of weighted distances to model noise can be seen as a special case of a much more general way to integrate side information into a metric. Several claims in the paper about these weighted distances such as being Lipschitz or their relationship to the unweighted distance are treated here in extreme generality.

Let (A, d_A) and (B, d_B) be metric spaces with $0 \in B$. Now let X be a compact subset of $A \times B$ and define

$$d_{A \times B}((a, b), (a', b')) = (d_A(a, a')^p + d_B(b, b')^p)^{1/p}$$

for $(a, b), (a', b') \in A \times B$. We will define the distance function $d_X : A \to \mathbb{R}$ to be the minimum distance from a point $a \in A$ to the set X.

$$d_X(a) := \min_{x \in X} d_{A \times B}((a, 0), x).$$

LEMMA A.1. (1-LIPSCHITZ) d_X *is 1-Lipschitz.*

Proof. Pick any $a, a' \in A$ and let $x' \in X$ be such that $d_X(a') = d_{A \times B}((a', 0), x')$. Note

$$\begin{aligned}
d_x(a) &= \min_{x \in X} d_{A \times B}((a, 0), x) \\
&\leq d_{A \times B}((a, 0), x') \\
&\leq d_{A \times B}((a, 0), (a', 0)) + d_{A \times B}((a', 0), x') \\
&= d_A(a, a') + d_X(a').
\end{aligned}$$

COROLLARY A.1. (d_k 1-LIPSCHITZ) *If $A \subseteq X$ is a weighted, compact set then for all $x, y \in X$*

$$d_k(x, A) - d_k(y, A) \leq d(x, y).$$

LEMMA A.2. (BOUNDED) $d_A(\,\cdot\,, X) \leq d_X$.

Proof. Let $x_0 \in X$ be such that $d_X(a) = d_{A \times B}((a, 0, x_0))$ and note,

$$\begin{aligned}
d_A(a, x) &= \min_{x \in X} d_A(a, x) \\
&\leq d_A(x, x0) \\
&\leq (d_A(a, x_0)^p + d_B(0, x_0)^p)^{1/p} \\
&= d_{A \times B}((a, 0), x_0) = d_X(a).
\end{aligned}$$

The following Lemma is used implicitly in the text to claim that if $A \subseteq B$ then $A^\delta \subset B^\delta$ for all $\delta \geq 0$.

LEMMA A.3. (MONOTONE) *If $B \subset X$ and $A \subset B$ are weighted compact sets then for all $x \in X$*

$$d_k(x, B) \leq d_k(x, A)$$

Proof. Note that for all $\epsilon > 0$ there exists an weighted k-set $Y_\epsilon \in \binom{A}{k}$ such that

$$\max_{y \in Y_\epsilon} \sqrt{\rho_y(q)} - \epsilon \leq d_k(q, A) \leq \max_{y \in Y_\epsilon} \sqrt{\rho_y(q)}$$

As $Y_\epsilon \in \binom{B}{k}$ if there exist any $b \in B \backslash A$ such that $\rho_b(x) < \max_{y \in Y_\epsilon} \rho_y(x)$ then $d_k(x, B) < d_k(x, A)$. Otherwise $d_k(x, B) = d_k(x, A)$.

B   Persistent Multi-cover Nerves

In this appendix, we discuss the multi-cover version of the Persistent Nerve Lemma (our Lemma 2.1 above). Although a combinatorial proof of this fact appeared in [16], we present here a summary of a more algebraic proof. In fact, the description below is virtually identical to the prove of the Nerve Theorem as found in Koslov [12], witht eh change that we are considering only a subcomplex of the barycentric subdivision. The interested reader can compare this to the standard setting (1-coverage) to see that the technical lemmas are all identical.

Let $\Lambda_k = \{v \subseteq [n] \mid |v| \geq k\}$ and for all $w \subseteq v \subseteq [n]$, we have a morphism $v \to w$. This is a subcategory of the so-called cosimplicial indexing category $\Delta^{\mathrm{op}} = \Lambda_0$. Given a collection $U = \{U_1, \ldots, U_n\}$ of open sets and $v \in \Lambda_k$, let $D(v) := \bigcap_{i \in v} U_i$. We say U is a *good open cover* if $D(v)$ is empty or contractible for all $v \in \Lambda_k$. For each morphism $v \to w$, let $D(v \to w)$ be the inclusion $D(v) \hookrightarrow D(w)$. The functor $D : \Lambda_k \to \mathrm{Top}$ is called a *diagram.* The homotopy colimit of D is

$$\mathrm{hocolim}\, D := \bigcup_{\sigma = v_0 \to \cdots \to v_m} (\underline{\sigma} \times D(v_0))$$

In this definition, $\sigma = v_0 \to \cdots \to v_m$ is a chain of morphisms in Λ_k and $\underline{\sigma}$ denotes the geometric realization of σ defined as

$$\underline{\sigma} := \{(x_v)_{v \in \Lambda_k} \mid x_v \in [0, 1], \sum_{v \in \Lambda_k} x_v = 1, x_v \neq 0 \text{ only if } v \in \sigma\}.$$

We use the geometric realization because when all maps in the diagram are inclusions, it allows us to avoid writing the homotopy colimit as a quotient space. This is closely related the blowup complex used in work on localized homology [19] and parallel algorithms for persistent homology [13].

Let $W_k U = \bigcup_{v \in \Lambda_k} D(v)$ denote the *k-covered region.* (This is also the colimit of D). Let $N_k U$ be the *k-nerve* of the cover U defined as the simplicial complex whose simplices are those chains of morphisms $\sigma = v_0 \to \cdots \to v_m$ such that all $D(v_i)$ are nonempty. It can be viewed as the subcomplex of the barycentric subdivision of the nerve of U induced on those vertices corresponding to simplices of dimension at least $k - 1$. For this reason, $N_0 U$ is homeomorphic to the nerve of U. More importantly, $N_k U$ is also equivalent to a diagram over Λ_k in which every space is a single point.

There are two natural maps associated with hocolim D. The first is the *base projection map* $b :$ hocolim $D \to N_k U$. The second is the *fiber projection map* $f :$ hocolim $D \to W_k U$. These maps are just the projections onto the first and second factors respectively of the product structure in hocolim D. As such, the projections commute with inclusions such as when we have

a pair of nested open covers. When U is a good open cover, the base projection map is a homotopy equivalence. This follows from the homotopy lemma for homotopy colimits, which says that a map between homotopy colimits where the maps on individual spaces are homotopy equivalences is itself a homotopy equivalence (see Koslov [12]). For open covers, the fiber projection map is also a homotopy equivalence. The main fact needed is that an open cover admits a partition of unity subordinate to it that allows one to define a lifting of W_k into a deformation retract of hocolim D. The proof of this fact may also be found in Koslov [12].

Let $V = \{V_1, \ldots, V_n\}$ be a good open subcover of U, that is $V_i \subseteq U_i$ for all i (note that some V_i could be empty). It follows that $W_k V \subseteq W_k U$ and also that $N_k V \subseteq N_k U$. Moreover, if D' is the diagram of V, then hocolim $D' \subseteq$ hocolim D and both projection maps commute with this inclusion. Thus, the homotopy equivalences $W_k U \to N_k U$ and $W_k V \to N_k V$ give a homotopy equivalence of pairs $(W_k U, W_k V) \to (N_k U, N_k V)$. For the same reason, if we have two pairs of good open covers $(U, V) \subseteq (U', V')$, then the isomorphisms $(W_k U, W_k V) \to (N_k U, N_k V)$ and $(W_k U', W_k V') \to (N_k U', N_k V')$ commute with the inclusions at the homotopy level (and thus also in homology). This last fact is sometimes known as the Persistent Nerve Lemma in the special case of $k = 0$ [3].